Real-Time Fraud Detection at Scale: An Advanced Architecture for Combating E-Commerce Review Manipulation

Combating e-commerce review manipulation with advanced architecture and AI techniques.

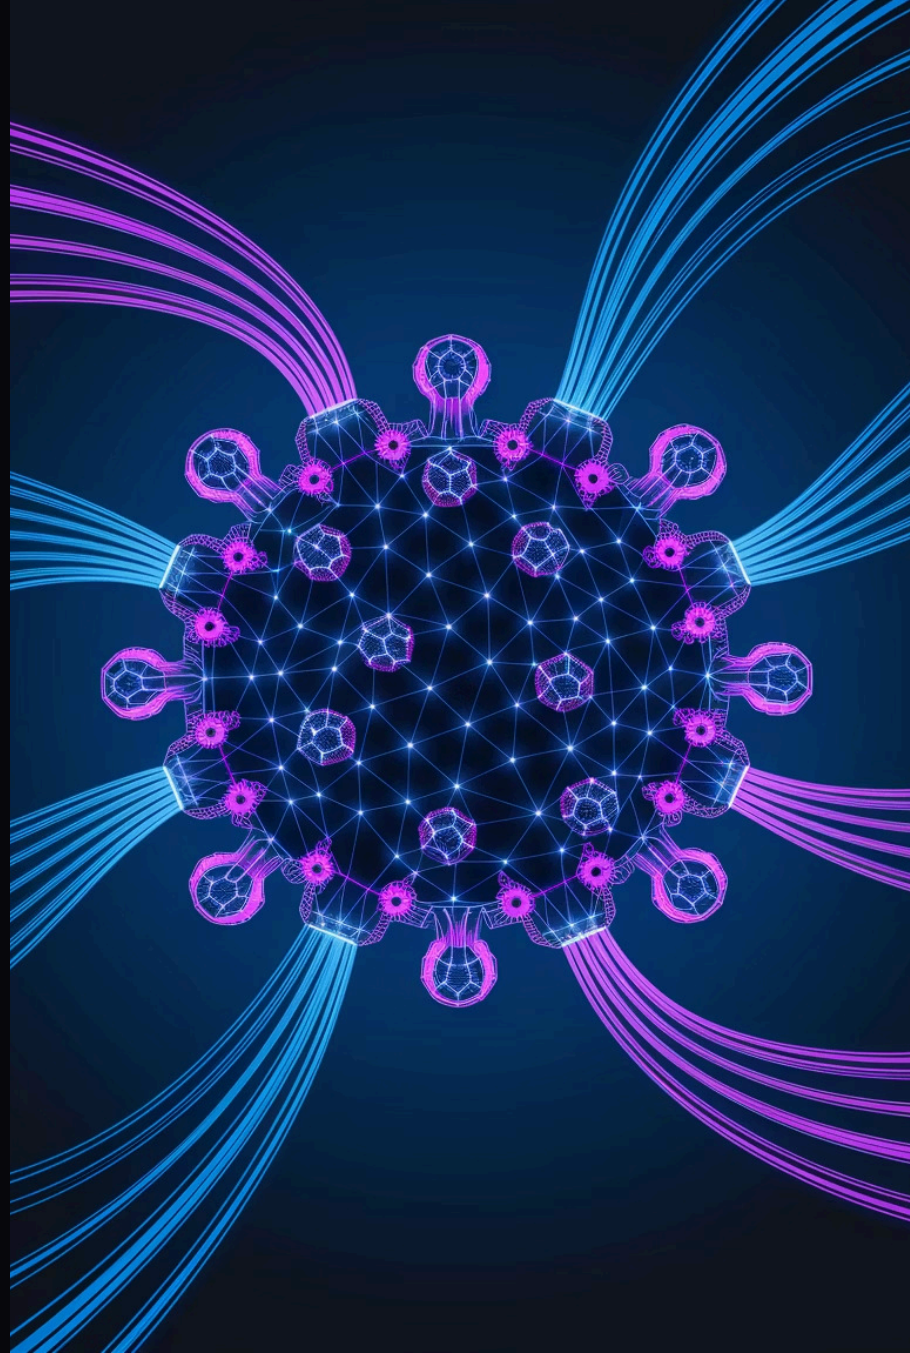

Satyanandam Kotha

JawaharLal Nehru Technological University

The Marketplace Integrity Challenge

Trust Erosion

Fake reviews undermine consumer confidence in platforms.

Economic Impact

Manipulated reviews distort purchasing decisions.

Sophisticated Fraud

Coordinated campaigns with multiple accounts evade detection.

Speed Requirements

Detection must happen before damage spreads.

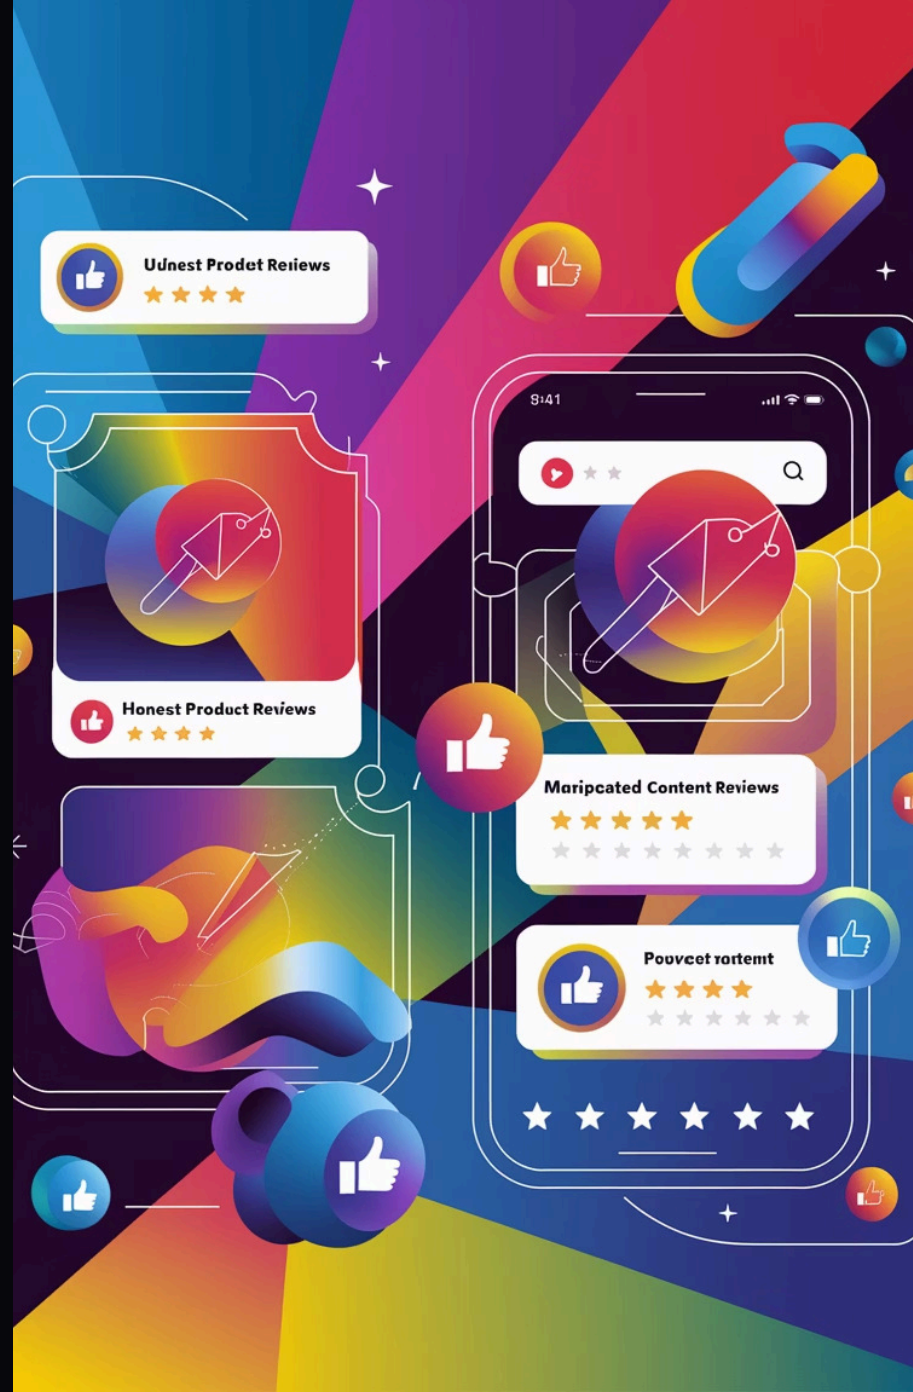

System Architecture Overview

Streaming Ingestion

Apache Kafka ingests and queues millions of reviews per second with fault-tolerant distribution.

Graph Storage

Custom graph database optimized for traversing complex review networks containing billions of edges.

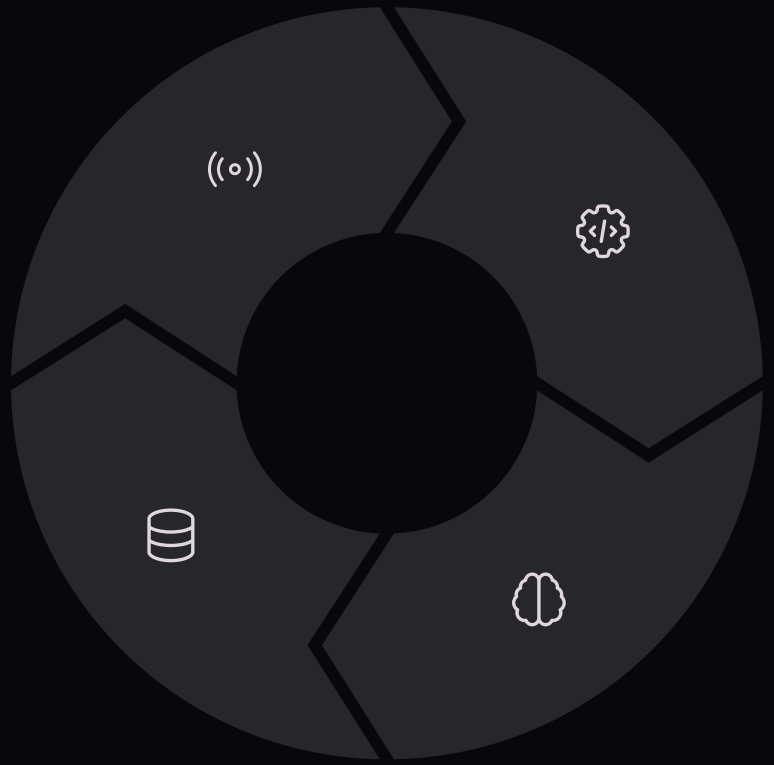

Real-Time Processing

Apache Flink enables stateful stream processing with consistent sub-100ms latency guarantees.

ML Inference

Horizontally scalable model serving infrastructure delivers fraud predictions within milliseconds.

Hybrid Detection Approach

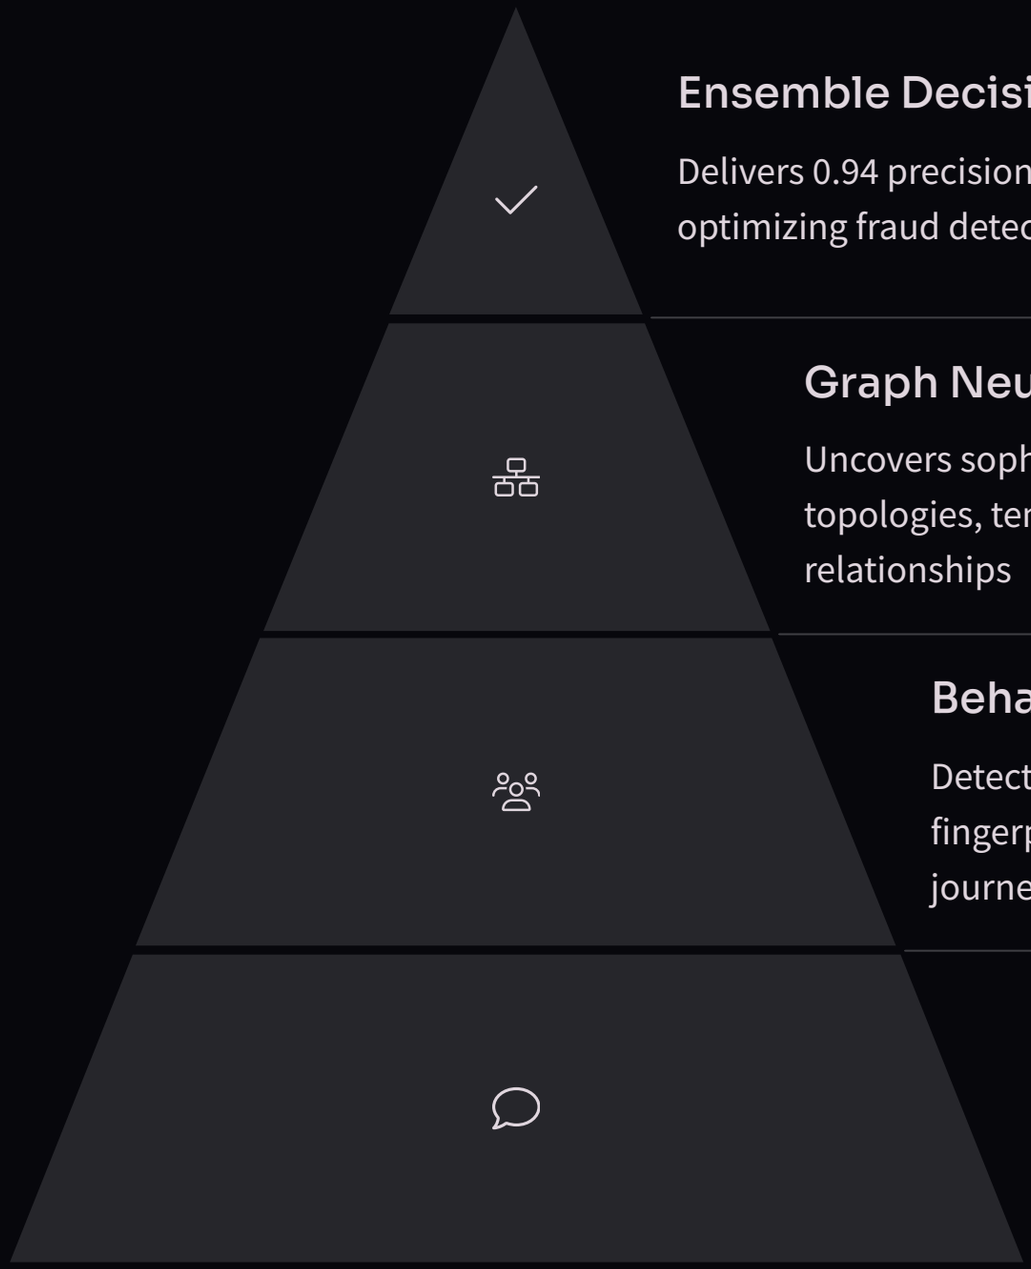

Ensemble Decision

Delivers 0.94 precision through strategic weighted voting across specialized models, optimizing fraud detection while minimizing false positives

Graph Neural Networks

Uncovers sophisticated coordination by analyzing complex network topologies, temporal submission patterns, and multi-dimensional account relationships

Behavioral Analytics

Detects anomalies by examining velocity patterns, device fingerprints, and sequential interaction signatures across the user journey

Natural Language Processing

Differentiates authentic content from AI-generated text using transformer-based semantic analysis, stylometric measurements, and cross-review consistency verification

Performance Metrics

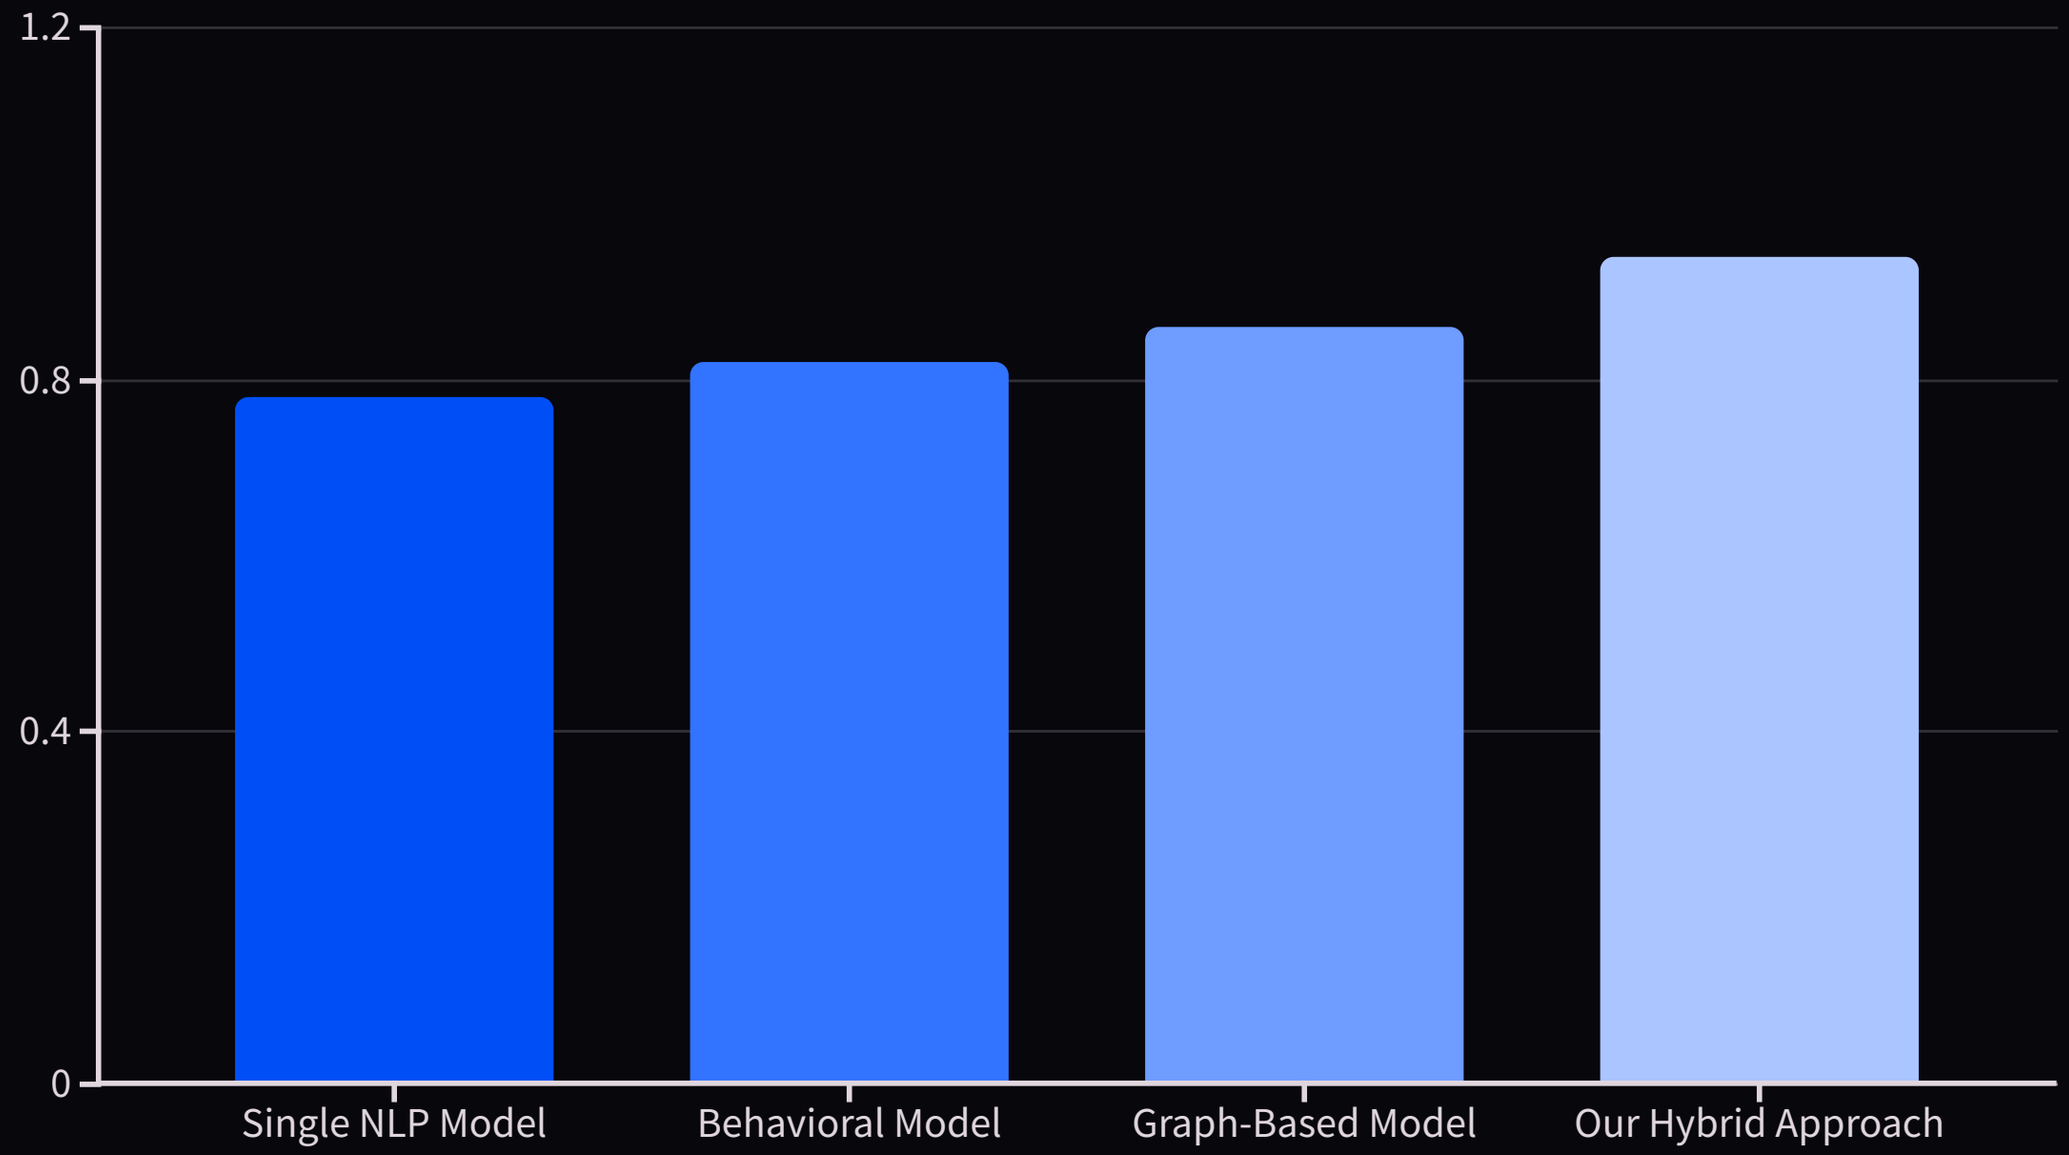

Our multi-model integration significantly outperforms single-architecture solutions across all key metrics.

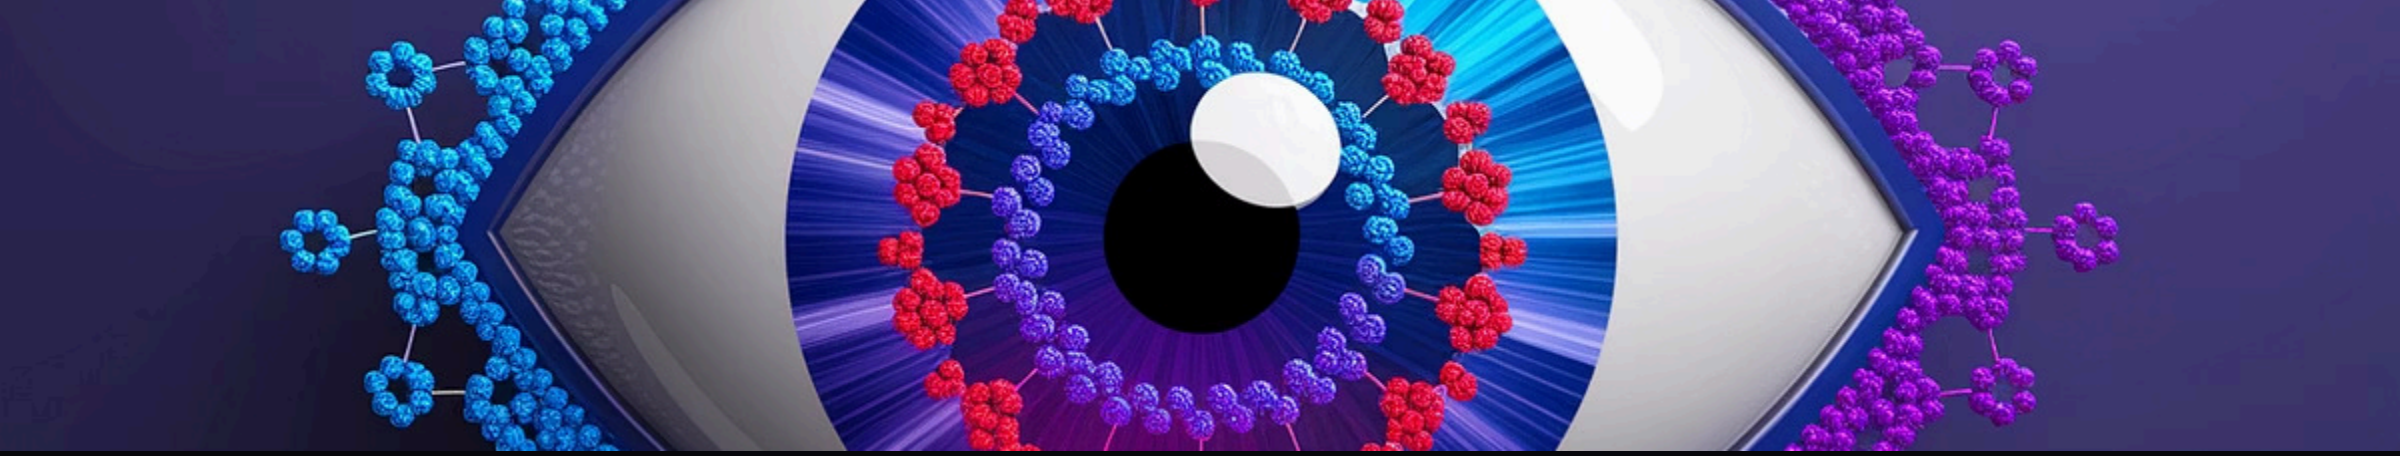

Coordinated Campaign Detection

8

Minimum Accounts

Threshold for identifying orchestrated review manipulation networks

12

Collusion Detection

Critical mass of accounts required for graph relationship analysis

3B

Edge Capacity

Maximum relationship connections our system can process in real-time

Our advanced temporal pattern analysis enables fraud detection at earlier stages with significantly fewer accounts than traditional methods, increasing prevention effectiveness.

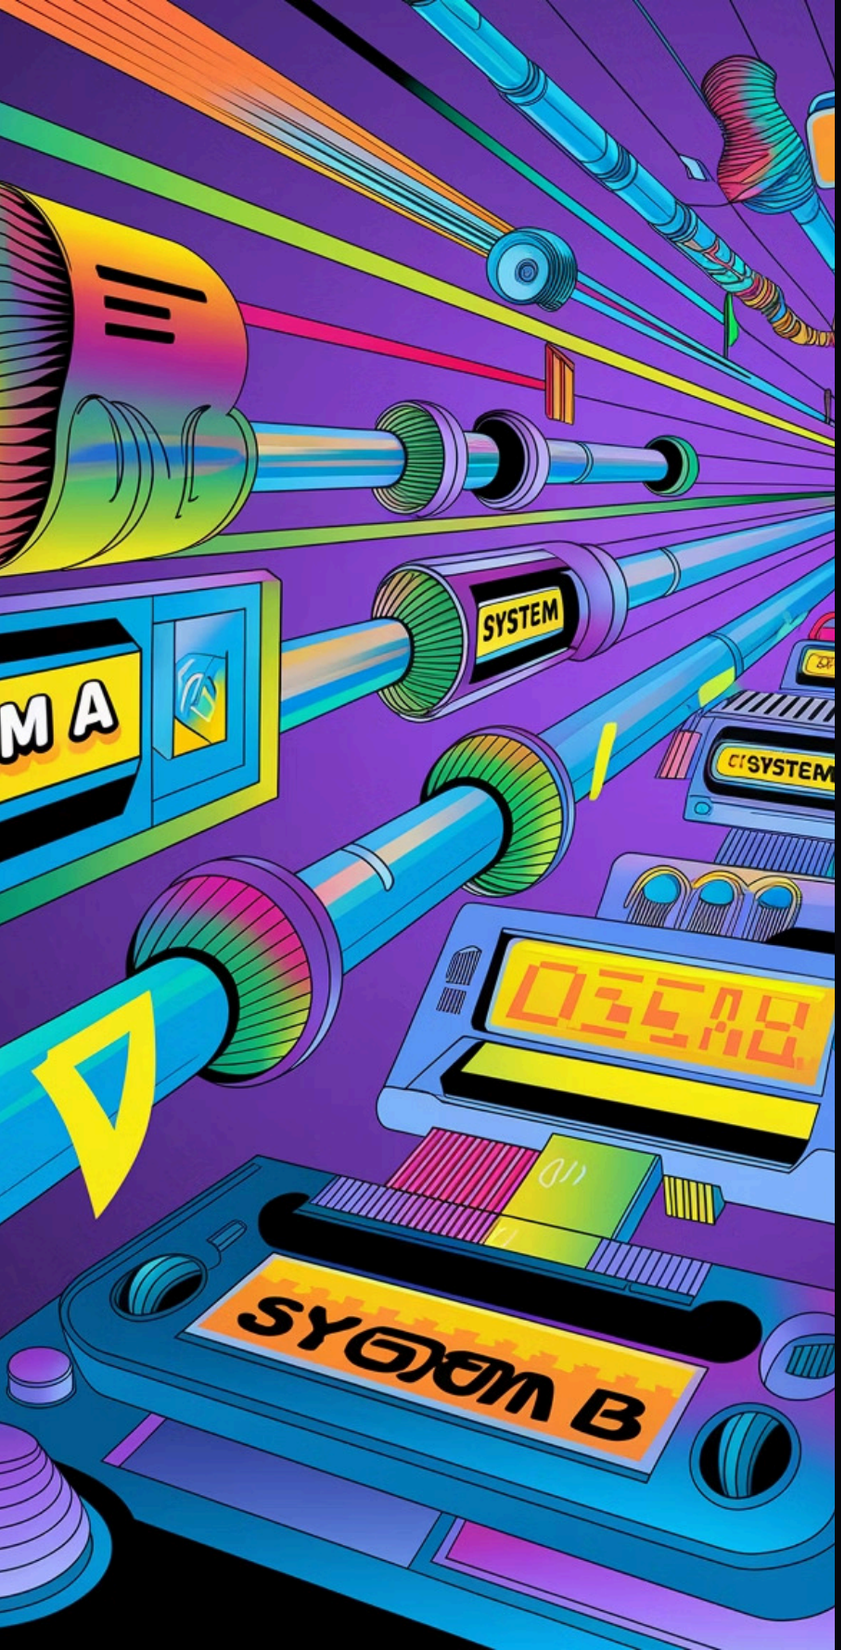

Privacy-Preserving Detection

Federated Learning

Models train collaboratively across multiple marketplace platforms without transferring raw user data.

Businesses maintain complete control over sensitive customer information within their own security perimeters.

Differential Privacy

Optimized epsilon value of 4.6 strikes the perfect balance between strong privacy guarantees and maintaining F1 detection scores above 0.92.

Mathematically proven safeguards prevent malicious actors from reverse-engineering individual user identities from aggregated data patterns.

Encrypted Inference

Cutting-edge homomorphic encryption enables sophisticated fraud analysis on fully encrypted datasets without decryption.

Delivers millisecond-level fraud detection while cryptographically shielding all personally identifiable information from exposure throughout the entire detection pipeline.

Scalability Performance

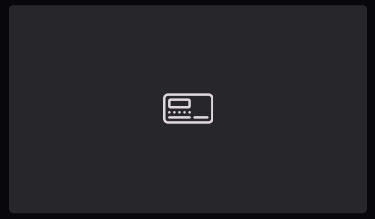

Near-Linear Scaling

Achieves 98% efficiency scaling to 32 worker nodes with only 2% performance overhead

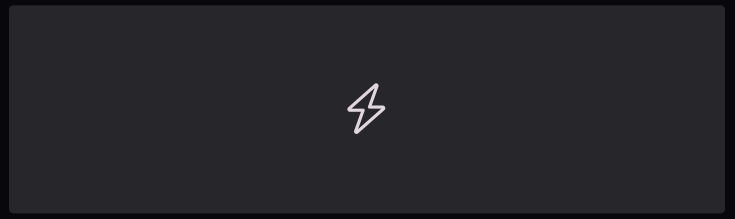

Sub-Linear Resource Usage

CPU and memory consumption grows at only 0.7x rate as traffic volume doubles

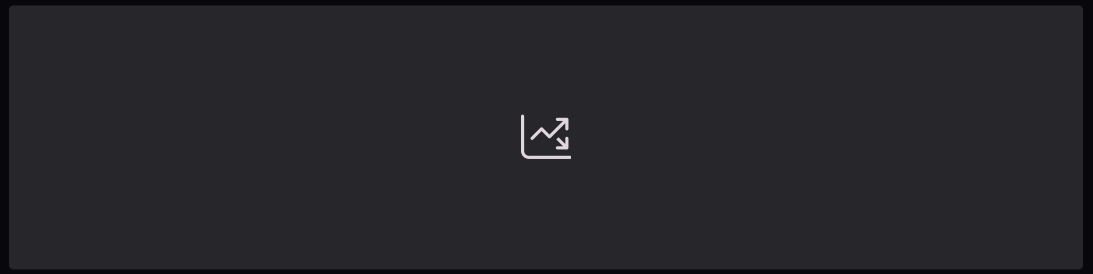

5x Traffic Handling

Maintains consistent 95ms response latency even during 500% traffic surge events

The architecture efficiently handles major promotional events when both legitimate and fraudulent review volumes increase by up to 10x, with automatic resource allocation ensuring uninterrupted detection capabilities.

Production Deployment

Platform Integration

Seamlessly integrated across 12+ major e-commerce ecosystems with consistent detection accuracy above 92%.

High Availability

Achieved 99.99% uptime through redundant infrastructure, geographic distribution, and automated failover mechanisms.

Fraud Reduction

Delivered 78% decrease in fraudulent reviews with 94% reduction in false positives across all market segments.

Continuous Improvement

Implemented automated daily model retraining pipeline that adapts to emerging fraud patterns within 24 hours of detection.

Implementation Insights

Design for Flexibility

Microservice-based architecture enables rapid integration of new detection algorithms without system downtime.

Data Structure Optimization

Custom sparse graph representations reduce memory footprint by 60% while improving query performance by 3x.

3 Stateful Processing

Leveraging Apache Flink's RocksDB state backends enables complex temporal fraud pattern recognition across user sessions.

Model Deployment Pipelines

CI/CD infrastructure automates A/B testing and canary deployments for seamless model transitions in production.

Key Takeaways

Speed Matters

Real-time detection under 100ms prevents fraudulent content from ever reaching users.

Multi-Model Approach

Ensemble architecture delivers 38% higher accuracy than any single detection model.

Privacy by Design

Federated learning preserves user confidentiality without compromising detection efficacy.

Scalable Architecture

System efficiently processes billions of connections with minimal latency during traffic spikes.

Thank you